

PATENT
5681-92600
SUN040309

"EXPRESS MAIL" MAILING LABEL NUMBER

EL990143000US

DATE OF DEPOSIT

4/8/04

I HEREBY CERTIFY THAT THIS PAPER OR FEE IS
BEING DEPOSITED WITH THE UNITED STATES
POSTAL SERVICE "EXPRESS MAIL POST OFFICE
TO ADDRESSEE" SERVICE UNDER 37 C.F.R. 1.10
ON THE DATE INDICATED ABOVE AND IS
ADDRESSED TO BOX PATENT APPLICATION,
COMMISSIONER FOR PATENTS, P.O. BOX 1450,
ALEXANDRIA, VA 22313-1450



Derrick Brown

METHOD AND APPARATUS

By:

Stuart A. Smithson
Martin Bate

Atty. Dkt. No.: 5681-92600

B. Noel Kivlin
Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Ph: (512) 853-8800

METHOD AND APPARATUS

Field of the Invention

[001] The present invention relates to a method of setting clock speeds for a computer system comprising a plurality of computing units. It is particularly applicable to a computer system comprising a plurality of computer processing units (CPUs), for example a plurality of computer servers in a single rack or a plurality of CPUs in the same server.

Background of the Invention

[002] Any computer system operating at a given clock frequency generates electromagnetic radiation of a frequency related to the clock frequency. Legislative and industry practice standards define maximum levels for these emissions so as to keep them within safe levels. Examples of formal tests of electromagnetic radiation in these situations are NEBs, FCC and European CE mark, and traditionally the intensity of radiation is measured by instruments with bandwidths of about 120KHz.

[003] When several CPUs are located in close proximity to each other, for example when multiple cores are involved, then their emissions tend to be cumulative and the additive effect of the emissions, particularly in a large concentration of CPUs, may inadvertently exceed legal standards at certain frequencies, even though emissions from each individual core meet the standard. Typically the higher emissions will be at harmonic frequencies or sub-harmonics of the CPU's clock and memory bus frequency (dependent on the number of CPUs).

[004] Previously this problem has been addressed by a spread spectrum technique in which the memory bus clock on a single server is continually varied across a predefined range. Such technology is built into a chip so that it is effectively integral with the CPU and it is thus relatively complex and cannot be retrofitted or easily altered or adapted.

Summary of the Invention

[005] The present invention provides a different clock frequency for each of a plurality of computing units of a computer system. Each different clock frequency is chosen to differ from

at least one other of said frequencies by at least a predetermined bandwidth.

[006] Preferably this minimum predetermined bandwidth is that which would allow the separation of frequencies to reduce the level of emissions within the measurement bandwidth to be below any given international legal electromagnetic emission limit. This separation could be from at least 1 Hz to at least 40 Ghz. Most conveniently it would be the bandwidth of a standard emission measuring device, which is typically 120KHz, and may for example advantageously be several times higher such as 1 MHZ, or 2 MHZ.

[007] In one embodiment the clock frequency at which each unit operates is determined as a function of a number generated, accessed or allocated on start-up or resetting of the system. For example such a number may be used as a multiplier of a base frequency which may be a difference between the clock frequency of two computing units.

[008] The number may be generated by one or more random number generators, or could be allocated by using a digit from the identifying address of the respective computing unit such as the CPU or server. Alternatively each different frequency could be dynamically allocated on start-up or resetting, using a start-up sequence which comprises identifying the core servers in a system, for example in a box or in a rack, but this method tends to be slower.

[009] The invention provides a relatively simple method which is generally straightforward to implement.

[010] The present invention also provides apparatus for start-up of a computer system comprising a plurality of computing units. The apparatus comprises means for setting a different clock frequency for each unit so that the units operate at a plurality of different frequencies. Each frequency differs from another of said frequencies by at least a predetermined minimum bandwidth. This frequency bandwidth may be a multiple of a predetermined base frequency. The multiple is determined by a respective number allocated on start-up or resetting of the system and the number may be generated by a random number generator or dynamically allocated.

Brief Description of Drawings

[011] For a better understanding of the present invention and to show how the same may be carried into effect, reference will now be made to examples of the invention, described with reference to the accompanying drawings in which:

Figure 1 is a schematic graph showing emission levels for a known computing system;

Figure 2 is a schematic graph showing emission levels for the computing system of the present invention;

Figure 3 illustrates a first embodiment computing system according to the present invention;

Figure 4 illustrates a second embodiment computing system according to the present invention;

Figure 5 is a flow diagram of the process of start-up of the apparatus of figures 1 and 2.

Detailed Description of Drawings

[012] The invention typically applies to a computer system such as the Sun Microsystems, Inc. Netra (TM) TI AC/DC 200 servers system in which six servers are mounted together in a single rack. Traditionally these servers would each be operated at the same CPU clock frequency of a value X between 600-700 MHz and preferably the higher speeds are used to optimize performance. The memory bus operates at one-sixth of the CPU speed, for example for a CPU speed of 648MHz the memory bus speed would be 108MHz. Emissions from such a system tend to reach peaks at multiple harmonics of the memory bus frequency, typically up to nine times the base frequency, ie up to 972MHz in this example. A typical legal limit for such emissions would be 57 dB μ V at 3m but the additive effect of the harmonics tends to exceed this at certain frequencies.

[013] Such an effect can be seen in the graph of Figure 1 which shows the cumulative effect on the level of radiation for a known multicore computer system, comprising up to six CPU systems 1 to 6 all running at frequency X. The legal limit which is typically 57 dB μ V at 3m, is indicated at A and it can be seen that the emission level exceeds this limit in this example when five or six of the CPU systems are operating. The cumulative effect is on a logarithmic scale ($20 \log (N \text{ total}/N \text{ original})$) and, for example, if each system emits at a level of 20 dB μ V

then the cumulative level of operating system 1 and system 2 would be 26 dB μ V, for systems 1 to 3 29.5 dB μ V, for systems 1 to 4 32 dB μ V, for five 14 dB μ V and for systems 1 to 6 would be 35.5 dB μ V. This is evidently undesirable.

[014] Figure 2 shows a typical level of radiation for a multicore computer system operated according to the invention. Here it can be seen that the limit shown at A is not exceeded, because each CPU system is operated at a different frequency separated by at least the bandwidth of the measuring device as indicated at 10. In figure 2 system 1 operates at frequency X, system 2 at X+n1, system 3 at X+n2, system 4 at X+n3, system 5 at X+n4 and system 6 at X+n5. Thus for the same emission level of 20 dB μ V for each CPU, the maximum level is 20 dB μ V regardless of the number of CPU systems used or added. One example of such system is illustrated in figure 3 where six systems 1 to 6 are mounted in a rack 11. Each system or core 1 to 6 has a respective CPU 31 to 36 operating at a different frequency. System 1 has a CPU 21 operating at frequency X MHz; system 2 has a CPU 22 operating at frequency (X + 1) MHz; system 3 has a CPU 23 operating at frequency (X + 2) MHz; system 4 has a CPU 24 operating at frequency (X + 3) MHz; system 5 has a CPU 25 operating at frequency (X + 4) MHz; and system 6 has a CPU 26 operating at frequency (X + 5) MHz.

[015] Figure 4 is a highly schematic view of an alternative embodiment of a computer system to which the invention can be applied. Six CPUs are shown schematically at 1, 2, 3, 4, 5 and 6 mounted in box II and connected by system bus 20. A clock controller for each CPU is shown schematically at 31, 32, 33, 34, 35 and 36 respectively together with a random number generator for each CPU respectively labelled 41, 42, 43, 44, 45 and 46. Of course, this is only one example: a single random number generator may be used for all CPUs or other means of allocating or generating numbers.

[016] On start-up (or resetting) of the computer system, each random number generator 41 - 46 generates a number which is used by the respective clock controller 31 - 36 as a multiplier to set the operating frequency of the respective CPU 1 to 6. Typically each random number generator 41 - 46 is set to generate a number between 1 and 6 corresponding to the total number at CPUs in the system and the clock controllers 31 - 36 multiply this number by a predetermined frequency difference figure of at least 120KHz and preferably 1 or 2 MHz to set the frequency difference between each of the operating frequencies of the CPUs.

[017] In this example the frequency difference figure n is 4 MHZ and the CPUs 1 to 6 operate at respective frequencies 648, 644, 640, 636, 632 and 628 MHZ.

[018] Figure 5 is a flow diagram illustrating a start-up procedure according to the invention for a multicore computer system such as that shown in figures 3 and 4. On start-up 71 a first frequency X is selected 72 and applied 73 to a first CPU of the system, such as CPU 61 in figure 4. The procedure then moves to the next CPU 74. A number N is then generated 75 either by selection from a set of stored numbers or by a random number generator such as 41 in figure 3, or by selecting a digit from the address of a server in the system. The number N is then used to select 75 a second frequency $Y=X+N$ and this frequency Y is applied 77 to the next CPU. This cycle is repeated until all CPUs or servers in the system have clock frequencies allocated to them.